Image and Multimedia Mining

Christos Faloutsos
CMU

Outline

- <u>Problem definition and motivation</u>
- Indexing and feature extraction
- Features for images (DWT, morphology, etc)
- Tools and case studies: PCA, ICA, Random Walks
- Conclusions

Problem #1: Similarity search

- [Bob Murphy] Sub-cellular protein localization patterns

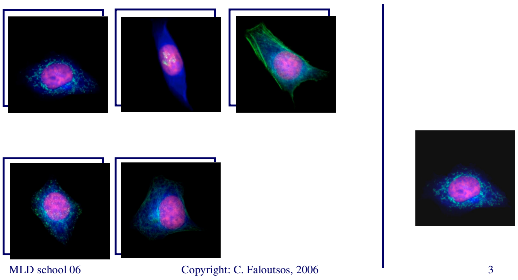

Problem #1: Similarity search

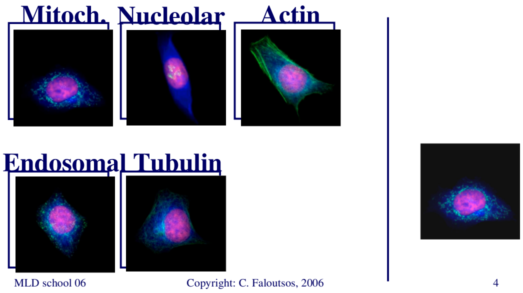

Problem#1': Similarity search:

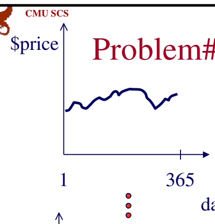

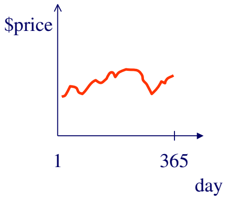

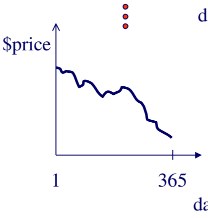

distance function? <later>

Problem #2: Mining

- [Bob Murphy] Sub-cellular protein localization patterns

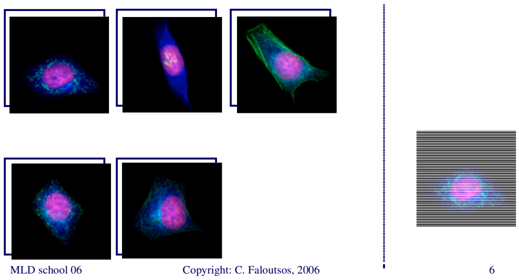

Problem#2’: Mining

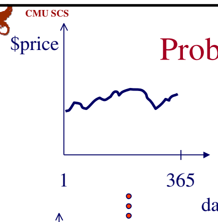

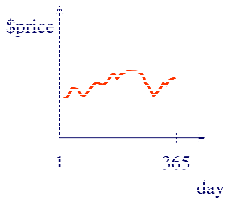

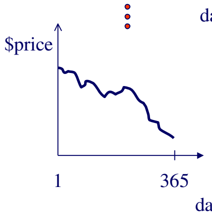

distance function? <later>

Problem definition

- Given: many sequences (images, video-clips)
  $x_1 , x_2 , \ldots , x_t , \ldots$
  $(y_1, y_2, \ldots , y_t, \ldots$
  $\ldots)$
- Find
  - similar sequences / images / video-clips
  - patterns; clusters; outliers

Motivation - applications: Images

- medical/biological images (training; research)
- biometrics/security
- satellite image analysis
- photo collections
- museum images
- logos
- …

Motivation - applications: Video

- surveillance
- TV news processing (eg., summarization - www.informedia.cs.cmu.edu)
- 3-d images (medical / biological)
- 3-d and 4-d datasets (x,y,z, time, temp., humidity etc) - weather/environment monitoring

Motivation - applications: Time sequences

- Financial, sales, economic series
- Medical (ECGs/EKGs, monitoring)
- civil infrastructure; automobile traffic monitoring

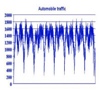

Motivation

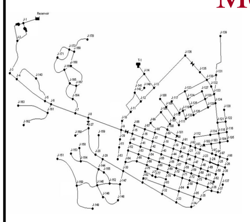

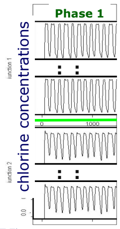

water distribution network
(w/ Jeanne Vanbriesen, CMU)

Motivation - Applications (cont'd)

- Weather, environment/anti-pollution
- computer network traffic monitoring
- data-center traffic monitoring (*self-** system at CMU)

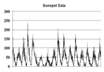

Outline

- Problem definition and motivation
- Indexing and feature extraction
- Features for images (DWT, morphology, etc)
- Tools and case studies: PCA, ICA, Random Walks
- Conclusions

Indexing

Problem:

- given a set of time sequences,
- find the ones similar to a desirable query sequence

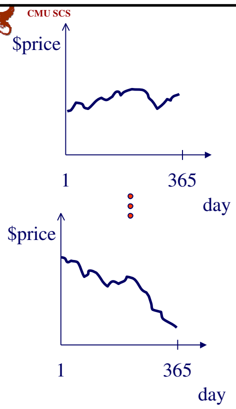

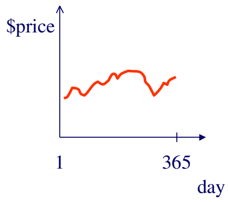

distance function: by expert

Idea: 'GEMINI'

Eg., '*find stocks similar to MSFT*'

Seq. scanning: too slow

How to accelerate the search?

[Faloutsos96]

'GEMINI' - Pictorially

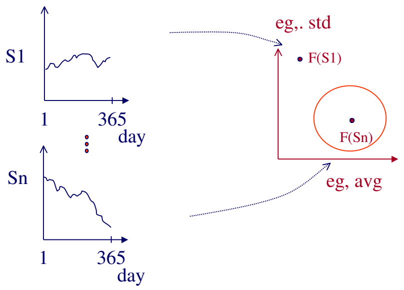

CMU SCS

GEMINI

Solution: Quick-and-dirty' filter:

- extract *n* features (numbers, eg., avg., etc.)
- map into a point in *n*-d feature space
- organize points with off-the-shelf spatial access method ('SAM')
- discard false alarms

CMU SCS

Examples of GEMINI

- Time sequences: DFT (up to 100 times faster) [SIGMOD94];
- [Kanellakis+], [Mendelzon+]

CMU SCS

Examples of GEMINI

Also, on many other data types:

- Images (QBIC) [JIIS94]
- tumor-like shapes [VLDB96]
- video [Informedia + S-R-trees]
- automobile part shapes [Kriegel+97]

CMU SCS

Indexing - SAMs

Q: How do Spatial Access Methods (SAMs) work?

A: they group nearby points (or regions) together, on nearby disk pages, and answer spatial queries quickly ('range queries', 'nearest neighbor' queries etc)

For example:

CMU SCS

R-trees

adv

- [Guttman84] eg., w/ fanout 4: group nearby rectangles to parent MBRs; each group -> disk page

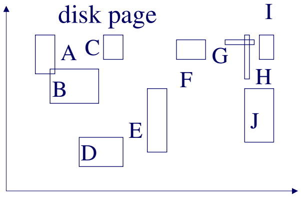

CMU SCS

R-trees

adv

- eg., w/ fanout 4:

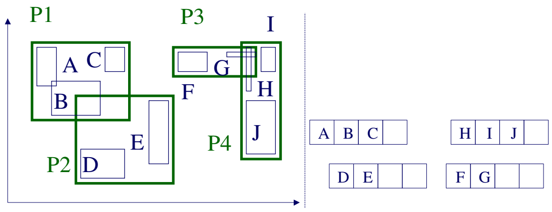

ML

adv

R-trees

- eg., w/ fanout 4:

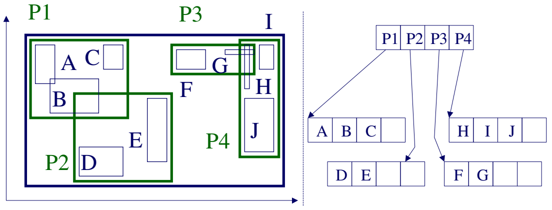

ML

adv

R-trees - range search?

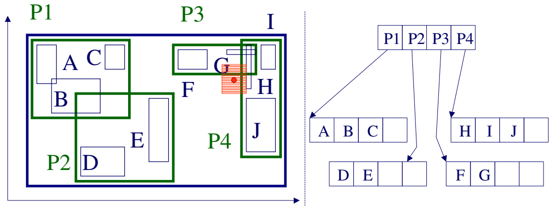

ML

adv

R-trees - range search?

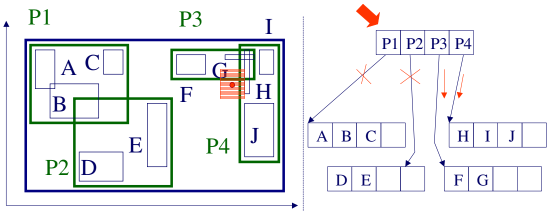

ML

Conclusions

- Fast indexing: through GEMINI
  - feature extraction and
  - (off the shelf) Spatial Access Methods [Gaede+98]

ML

Outline

- Problem definition and motivation
- Indexing and feature extraction
- Features for images (DWT, morphology, etc)
- Tools and case studies: PCA, ICA, Random Walks
- Conclusions

ML

Outline

- Problem definition and motivation
- Indexing and feature extraction
  - DFT, DWT etc
  - SVD/PCA
  - MDS and FastMap
- ...

DFT and cousins

- very good for compressing real signals
- more details on DWT: later

DFT and stocks

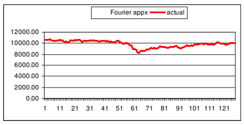

- Dow Jones Industrial index, 6/18/2001-12/21/2001

DFT and stocks

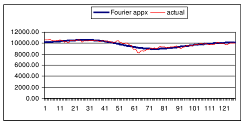

- Dow Jones Industrial index, 6/18/2001-12/21/2001
- just 3 DFT coefficients give very good approximation

Log(ampl)

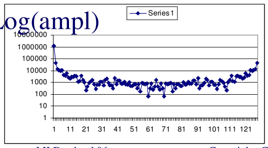

freq

Outline

- Problem definition and motivation
- Indexing and feature extraction
  - DFT, DWT etc
  - SVD/PCA
  - MDS and FastMap
- ...

SVD

- THE optimal method for dimensionality reduction
  - (under the Euclidean metric)

SVD - Motivation

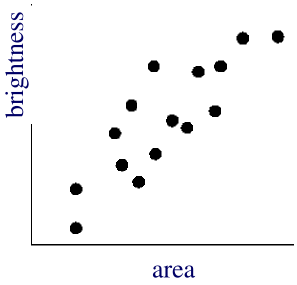

CMU SCS

SVD - Motivation

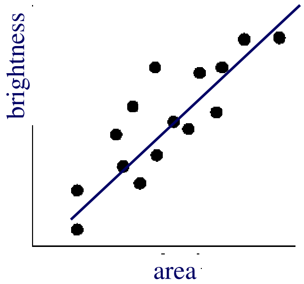

MLD school 06 Copyright: C. Faloutsos, 2006 37

CMU SCS

SVD – dim. reduction

SVD: gives
best axis to project

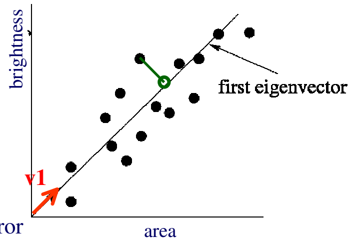

- minimum RMS error

MLD school 06 Copyright: C. Faloutsos, 2006 38

CMU SCS

Singular Value Decomposition (SVD)

- SVD (~LSI ~ KL ~ PCA ~ spectral analysis...)
- LSI: S. Dumais; M. Berry
- KL: eg, Duda+Hart
- PCA: eg., Jolliffe
- Details: [Press+], [Faloutsos96]

MLD school 06 Copyright: C. Faloutsos, 2006 39

CMU SCS

SVD

- **Extremely** useful tool
  - (also behind PageRank/google and Kleinberg's algorithm for hubs and authorities)
- But may be slow: O($N * M * M$) if $N>M$
- any approximate, faster method?

MLD school 06 Copyright: C. Faloutsos, 2006 40

CMU SCS

SVD shorcuts

- random projections (Johnson-Lindenstrauss thm [Papadimitriou+ pods98])

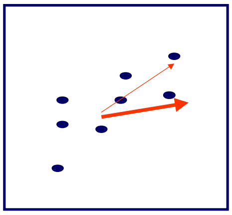

MLD school 06 Copyright: C. Faloutsos, 2006 41

CMU SCS

Random projections

- pick 'enough' random directions (will be ~orthogonal, in high-d!!)
- distances are preserved probabilistically, within epsilon
- (also, use as a pre-processing step for SVD [Papadimitriou+ PODS98])

MLD school 06 Copyright: C. Faloutsos, 2006 42

Outline

- Problem definition and motivation
- Indexing and feature extraction
  - DFT, DWT etc
  - SVD
  - MDS and FastMap
- ...

MDS / FastMap

- but, what if we have NO points to start with?
  (eg. Time-warping distance)
- A: Multi-dimensional Scaling (MDS) ; FastMap

MDS/FastMap

	O1	O2	O3	O4	O5
O1	0	1	1	100	100
O2	1	0	1	100	100
O3	1	1	0	100	100
O4	100	100	100	0	1
O5	100	100	100	1	0

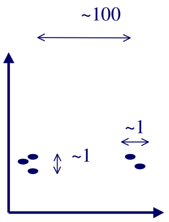

MDS

Multi Dimensional Scaling

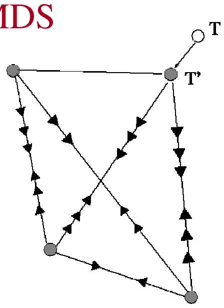

FastMap

- Multi-dimensional scaling (MDS) can do that, but in O(N**2) time
- FastMap [Faloutsos+95] takes O(N) time

FastMap: Application

VideoTrails [Kobla+97]

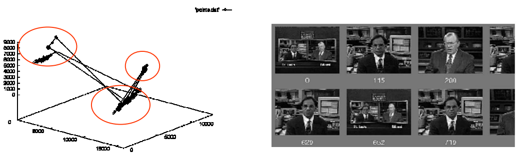

scene-cut detection (about 10% errors)

CMU SCS

Conclusions - Practitioner's guide

Similarity search in time sequences

1) establish/choose distance (Euclidean, time-warping,…)

2) extract features (SVD, DWT, MDS), and use an SAM (R-tree/variant) or a Metric Tree (M-tree)

2') for high intrinsic dimensionalities, consider sequential scan (it might win…)

CMU SCS

Books

- William H. Press, Saul A. Teukolsky, William T. Vetterling and Brian P. Flannery: *Numerical Recipes in C*, Cambridge University Press, 1992, 2nd Edition. (Great description, intuition and code for SVD)
- C. Faloutsos: *Searching Multimedia Databases by Content*, Kluwer Academic Press, 1996 (introduction to SVD, and GEMINI)

CMU SCS

References

- Agrawal, R., K.-I. Lin, et al. (Sept. 1995). Fast Similarity Search in the Presence of Noise, Scaling and Translation in Time-Series Databases. Proc. of VLDB, Zurich, Switzerland.
- Babu, S. and J. Widom (2001). "Continuous Queries over Data Streams." SIGMOD Record 30(3): 109-120.
- Breunig, M. M., H.-P. Kriegel, et al. (2000). LOF: Identifying Density-Based Local Outliers. SIGMOD Conference, Dallas, TX.
- Berry, Michael: http://www.cs.utk.edu/~lsi/

CMU SCS

References

- Ciaccia, P., M. Patella, et al. (1997). M-tree: An Efficient Access Method for Similarity Search in Metric Spaces. VLDB.
- Foltz, P. W. and S. T. Dumais (Dec. 1992). "Personalized Information Delivery: An Analysis of Information Filtering Methods." Comm. of ACM (CACM) 35(12): 51-60.
- Guttman, A. (June 1984). R-Trees: A Dynamic Index Structure for Spatial Searching. Proc. ACM SIGMOD, Boston, Mass.

CMU SCS

References

- Gaede, V. and O. Guenther (1998). "Multidimensional Access Methods." Computing Surveys 30(2): 170-231.
- Gehrke, J. E., F. Korn, et al. (May 2001). On Computing Correlated Aggregates Over Continual Data Streams. ACM Sigmod, Santa Barbara, California.

CMU SCS

References

- Gunopulos, D. and G. Das (2001). Time Series Similarity Measures and Time Series Indexing. SIGMOD Conference, Santa Barbara, CA.
- Hatonen, K., M. Klemettinen, et al. (1996). Knowledge Discovery from Telecommunication Network Alarm Databases. ICDE, New Orleans, Louisiana.
- Jolliffe, I. T. (1986). Principal Component Analysis, Springer Verlag.

References

- Keogh, E. J., K. Chakrabarti, et al. (2001). Locally Adaptive Dimensionality Reduction for Indexing Large Time Series Databases. SIGMOD Conference, Santa Barbara, CA.
- Kobla, V., D. S. Doermann, et al. (Nov. 1997). VideoTrails: Representing and Visualizing Structure in Video Sequences. ACM Multimedia 97, Seattle, WA.

References

- Oppenheim, I. J., A. Jain, et al. (March 2002). A MEMS Ultrasonic Transducer for Resident Monitoring of Steel Structures. SPIE Smart Structures Conference SS05, San Diego.
- Papadimitriou, C. H., P. Raghavan, et al. (1998). Latent Semantic Indexing: A Probabilistic Analysis. PODS, Seattle, WA.
- Rabiner, L. and B.-H. Juang (1993). Fundamentals of Speech Recognition, Prentice Hall.

References

- Traina, C., A. Traina, et al. (October 2000). Fast feature selection using the fractal dimension,. XV Brazilian Symposium on Databases (SBBD), Paraiba, Brazil.

References

- Dennis Shasha and Yunyue Zhu *High Performance Discovery in Time Series: Techniques and Case Studies* Springer 2004
- Yunyue Zhu, Dennis Shasha ``*StatStream: Statistical Monitoring of Thousands of Data Streams in Real Time*'' VLDB, August, 2002. pp. 358-369.
- Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. *The Design of an Acquisitional Query Processor for Sensor Networks*. SIGMOD, June 2003, San Diego, CA.

Outline

- Problem definition and motivation
- Indexing and feature extraction
- Features for images (DWT, morphology, etc)
- Tools and case studies: PCA, ICA, Random Walks
- Conclusions

How to extract features?

- Eg.: [Bob Murphy]: protein localization images

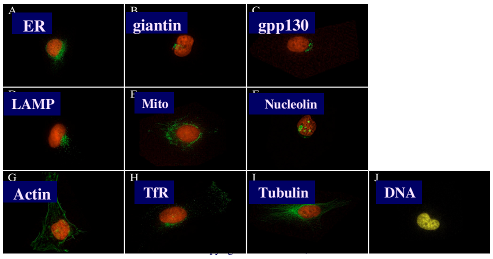

CMU SCS

How to extract features?

- Eg.:

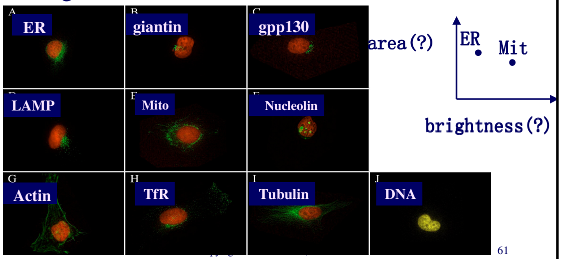

CMU SCS

(Natural) image features

[Flickner+, 95] from IBM's QBIC:

- Color
- Shape
- Texture

CMU SCS

(Natural) image features

[Flickner+, 95] from IBM's QBIC:

- Color: average R/G/B; color histograms
- Shape
- Texture

CMU SCS

Images - shapes

- distance function: Euclidean, on the area, perimeter, and 20 'moments' [QBIC, '95]
- Q: other 'features' / distance functions?

CMU SCS

Images - shapes

- (A1: turning angle)
- A2: wavelets
- A3: morphology: dilations/erosions
- ...

CMU SCS

Images - shapes

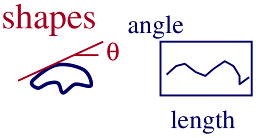

- (A1: turning angle)
- A2: wavelets
- A3: morphology: dilations/erosions
- ...

Wavelets - example

http://grail.cs.washington.edu/projects/query/

Wavelets achieve *great* compression:

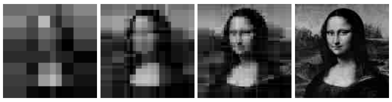

20 100 400 16,000

coefficients

Wavelets - intuition

- Edges (horizontal; vertical; diagonal)

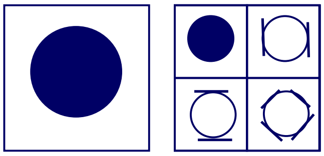

Wavelets - intuition

- Edges (horizontal; vertical; diagonal)
- recurse

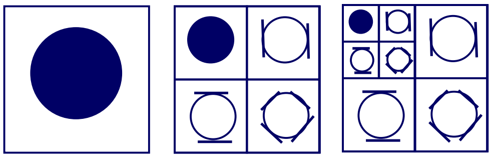

Wavelets - intuition

- Edges (horizontal; vertical; diagonal)
- http://www331.jpl.nasa.gov/public/wave.html

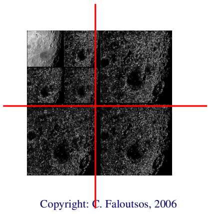

Wavelets

- Many wavelet basis:
  - Haar
  - Daubechies (-4, -6, -20)
  - Gabor
  - ...

Gabor Function

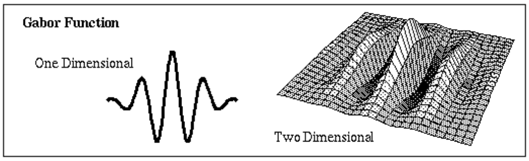

We can extend the function to generate Gabor filters by rotating and dilating

Wavelets:

- Extremely useful
- Excellent compression / feature extraction, for natural images
- fast to compute (*O(N)*)

Crush intro to Haar wavelets

- The simplest to code, explain and understand

Wavelets - DWT

- Similarly, DFT suffers on short-duration waves (eg., baritone, silence, soprano)

value

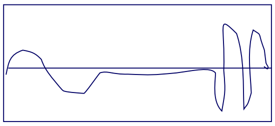

time

Haar Wavelets

- subtract sum of left half from right half
- repeat recursively for quarters, eight-ths, ...

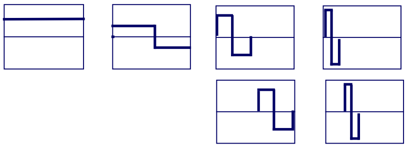

Wavelets - construction

adv

x0 x1 x2 x3 x4 x5 x6 x7

Wavelets - construction

adv

level 1 d1,0 s1,0 d1,1 s1,1 ....

\- +

x0 x1 x2 x3 x4 x5 x6 x7

CMU SCS

ML

Wavelets - construction

adv

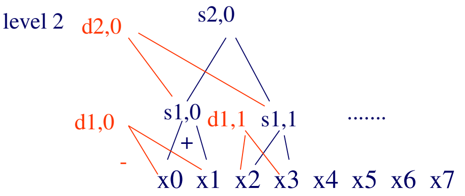

MLD school 06 Copyright: C. Faloutsos, 2006 79

CMU SCS

ML

Wavelets - construction

adv

etc ...

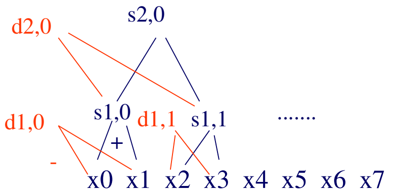

MLD school 06 Copyright: C. Faloutsos, 2006 80

CMU SCS

ML

Wavelets - construction

adv

Q: map each coefficient

on the time-freq. plane

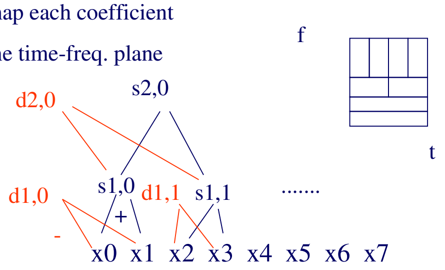

MLD school 06 Copyright: C. Faloutsos, 2006 81

CMU SCS

ML

Wavelets - construction

adv

Q: map each coefficient

on the time-freq. plane

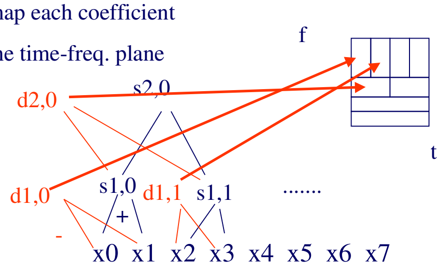

MLD school 06 Copyright: C. Faloutsos, 2006 82

CMU SCS

ML

Haar wavelets - code

adv

```
#!/usr/bin/perl5
# expects a file with numbers
# and prints the dwt transform
# The number of time-ticks should be a power of 2
# USAGE
#    haar.pl <fname>

my @vals=();
my @smooth; # the smooth component of the signal
my @diff;   # the high-freq. component

# collect the values into the array @val
while(<>){
        @vals = ( @vals ,  split );
}
```

```
my $len = scalar(@vals);
my $half = int($len/2);
while($half >= 1 ){
  for(my $i=0; $i< $half; $i++){
        $diff [$i] = ($vals[2*$i] - $vals[2*$i + 1] )/ sqrt(2);
        print "\t", $diff[$i];
        $smooth [$i] = ($vals[2*$i] + $vals[2*$i + 1] )/ sqrt(2);
  }
  print "\n";
  @vals = @smooth;
  $half = int($half/2);
}
print "\t", $vals[0], "\n" ;      # the final, smooth component
```

MLD school 06 Copyright: C. Faloutsos, 2006 83

CMU SCS

ML

Images - shapes

- (A1: turning angle)
- A2: wavelets
- A3: morphology: dilations/erosions
- ...

MLD school 06 Copyright: C. Faloutsos, 2006 84

Other shape features

- Morphology (dilations, erosions, openings, closings) [Korn+, VLDB96]

shape (**B/W**)

"structuring element"

Other shape features

- Morphology (dilations, erosions, openings, closings) [Korn+, VLDB96]

shape

"structuring element"

Other shape features

- Morphology (dilations, erosions, openings, closings) [Korn+, VLDB96]

shape

"structuring element"

Morphology: closing

- fill in small gaps
- **very similar** to 'alpha contours'

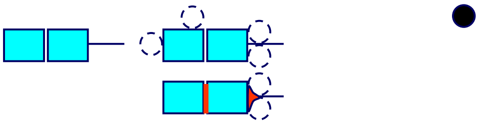

Morphology: closing

- fill in small gaps

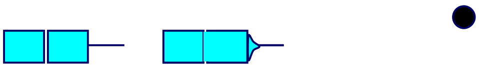

'closing', with R=1

Morphology: opening

- 'closing', for the complement =
- trim small extremities

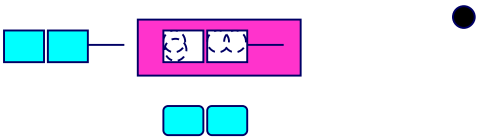

Morphology: opening

- ‘closing’, for the complement =
- trim small extremities

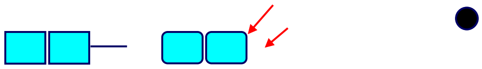

‘opening’
with R=1

Morphology

- Closing: fills in gaps

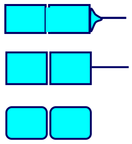

- Opening: trims extremities
- All wrt a structuring element:

Morphology

- Features: areas of openings (R=1, 2, …) and closings

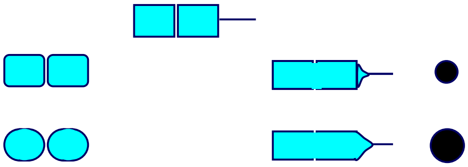

Morphology

- resulting areas: ‘pattern spectrum’
  - translation (and rotation) independent
- As described: on b/w images
  - can be extended to grayscale ones (eg., by thresholding)

Conclusions

Features for color, shape, texture

- Shape: wavelets; math. morphology
- texture: wavelets

we didn’t elaborate:

- color: color histograms; avg R/G/B
- and many more

References

- Faloutsos, C., R. Barber, et al. (July 1994). “*Efficient and Effective Querying by Image Content*.” J. of Intelligent Information Systems 3(3/4): 231-262.
- Faloutsos, C. and K.-I. D. Lin (May 1995). *FastMap: A Fast Algorithm for Indexing, Data-Mining and Visualization of Traditional and Multimedia Datasets*. Proc. of ACM-SIGMOD, San Jose, CA.
- Faloutsos, C., M. Ranganathan, et al. (May 25-27, 1994). *Fast Subsequence Matching in Time-Series Databases*. Proc. ACM SIGMOD, Minneapolis, MN.

CMU SCS

ML

References

- Christos Faloutsos, *Searching Multimedia Databases by Content,* Kluwer 1996

MLD school 06 Copyright: C. Faloutsos, 2006 97

CMU SCS

ML

References

- Flickner, M., H. Sawhney, et al. (Sept. 1995). "*Query by Image and Video Content: The QBIC System.*" IEEE Computer 28(9): 23-32.
- Goldin, D. Q. and P. C. Kanellakis (Sept. 19-22, 1995). *On Similarity Queries for Time-Series Data: Constraint Specification and Implementation* (CP95), Cassis, France.

MLD school 06 Copyright: C. Faloutsos, 2006 98

CMU SCS

ML

References

- Charles E. Jacobs, Adam Finkelstein, and David H. Salesin. *Fast Multiresolution Image Querying* SIGGRAPH '95, pages 277-286. ACM, New York, 1995.
- Flip Korn, Nikolaos Sidiropoulos, Christos Faloutsos, Eliot Siegel, Zenon Protopapas: *Fast Nearest Neighbor Search in Medical Image Databases*. VLDB 1996: 215-226

MLD school 06 Copyright: C. Faloutsos, 2006 99

CMU SCS

ML

Resources - software and urls

- http://www.dsptutor.freeuk.com/jsanalyser/FFTSpectrumAnalyser.html : Nice java applets for FFT
- http://www.relisoft.com/freeware/freq.html voice frequency analyzer (needs microphone)

MLD school 06 Copyright: C. Faloutsos, 2006 100

CMU SCS

ML

Resources: software and urls

- *xwpl:* open source wavelet package from Yale, with excellent GUI
- http://monet.me.ic.ac.uk/people/gavin/java/waveletDemos.html : wavelets and scalograms

MLD school 06 Copyright: C. Faloutsos, 2006 101

CMU SCS

ML

Books

- William H. Press, Saul A. Teukolsky, William T. Vetterling and Brian P. Flannery: *Numerical Recipes in C*, Cambridge University Press, 1992, 2nd Edition. (Great description, intuition and code for DFT, DWT)
- C. Faloutsos: *Searching Multimedia Databases by Content*, Kluwer Academic Press, 1996 (introduction to DFT, DWT)

MLD school 06 Copyright: C. Faloutsos, 2006 102

Outline

- Problem definition and motivation
- Indexing and feature extraction
- Features for images (DWT, morphology, etc)
- Tools and case studies: PCA, ICA, Random Walks
- Conclusions

ICA

- Independent Component Analysis
  - better than PCA/SVD
  - also known as ‘blind source separation’
  - (the `cocktail discussion’ problem)

Intuition behind ICA:

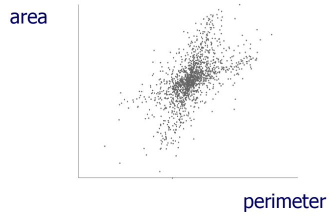

Intuition behind ICA:

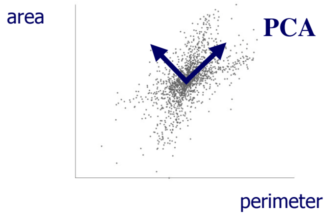

Intuition behind ICA:

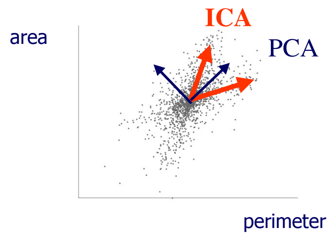

Intuition behind ICA:

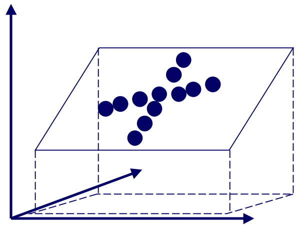

CMU SCS ML

Intuition behind ICA:

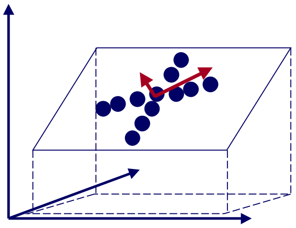

PCA finds the hyperplane.

CMU SCS ML

Intuition behind ICA:

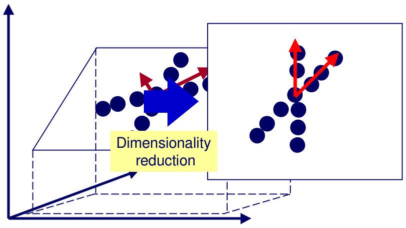

PCA finds the hyperplane. ICA finds the correct patterns.

CMU SCS ML

ICA in action

- stock prices - find groups / outliers!

CMU SCS ML

Eg.: Hidden variables in stock prices

DJIA (29 companies)
Unimodal data

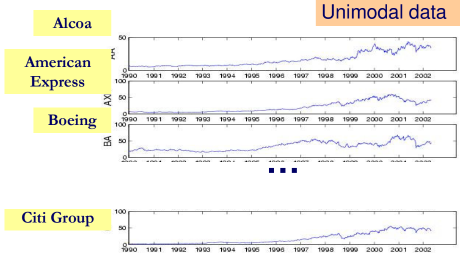

(Q) What patterns do you see?

CMU SCS ML

Eg.: Hidden variables in stock prices

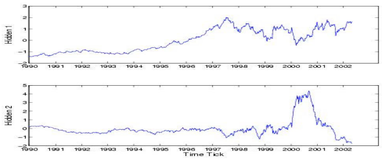

Discovery non-obvious hidden variables.
(How to do this automatically?)

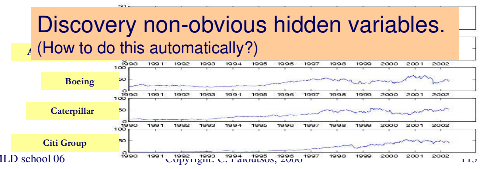

CMU SCS ML adv

Hidden variables

Dow Jones Industrial Average

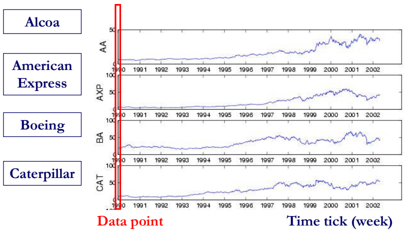

Goal:
Find meaningful hidden variables

Weekly DJIA closing prices (01/02/1990-08/05/2002)
n=660 data points, m=29 companies

CMU SCS

ML Math

(A.1) How did stock prices change?

29

AA1, …, XOM1 … … AAn, …, XOMn	=	H11, H12, …, H1m … ? … Hn1, Hn2, …, Hnm	B11, B12, …, B1m … ? Bm1, Bm2, …, Bmm

Time tick

n=660

Hidden variable

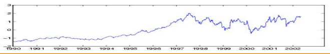

Time tick

CMU SCS

ML

Hidden variables by AutoSplit

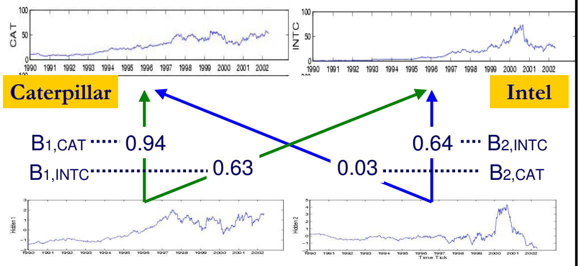

Hidden v#1 (meaning??) Hidden v#2 (meaning??)

CMU SCS

ML

Hidden variables by AutoSplit

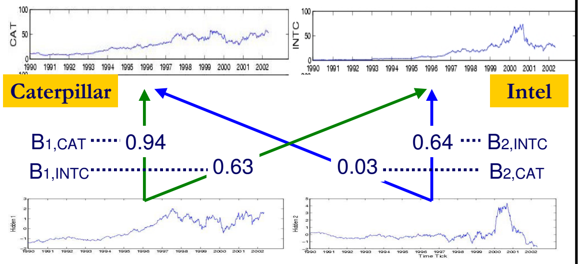

"General trend" "Internet bubble"

CMU SCS

ML

Making the most of ICA/AutoSplit

- Pattern discovery (hidden variables)
- Clusters (tech vs non-tech stocks)
- Outlier detection

CMU SCS

ML

Companies related to hidden variable 1

$B_{1,j}$			
Highest		Lowest	
Caterpillar	0.938512	AT&T	0.021885
Boeing	0.911120	WalMart	0.624570
MMM	0.906542	Intel	0.638010
Coca Cola	0.903858	Home Depot	0.647774
Du Pont	0.900317	Hewlett-Packard	0.658768

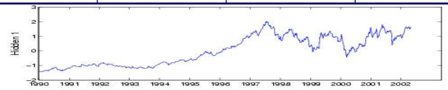

Hidden variable 1 : "General trend"

CMU SCS

ML

Outlier detection

"General trend"

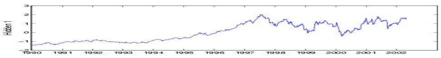

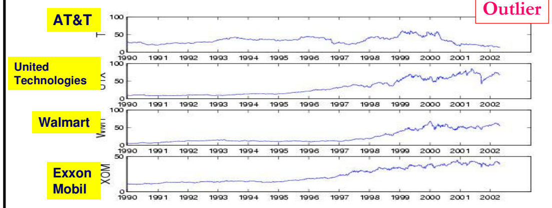

Conclusions for ICA

- Better than PCA
- Actually, uses PCA as a first step!

References

- Fukunaga, K. (1990). *Introduction to Statistical Pattern Recognition*, Academic Press.
- Jolliffe, I. T. (1986). *Principal Component Analysis*, Springer Verlag.
- Aapo Hyvarinen, Juha Karhunen, and Erkki Oja *Independent Component Analysis*, John Wiley & Sons, 2001.

References

- Jia-Yu Pan, Hiroyuki Kitagawa, Christos Faloutsos, and Masafumi Hamamoto. *AutoSplit: Fast and Scalable Discovery of Hidden Variables in Stream and Multimedia Databases*. PAKDD 2004

Outline

- ...
- Tools and case studies: PCA, ICA, Random Walks
  - ICA; case study: visual vocabulary (ViVo)
  - Random Walks
- Conclusions

ViVo: cat retina mining

- with Ambuj Singh, Mark Verardo, Vebjorn Ljosa, Arnab Bhattacharya (UCSB)
- Jia-Yu Tim Pan, HJ Yang (CMU)

Retina, its image, and the detachment

- retina

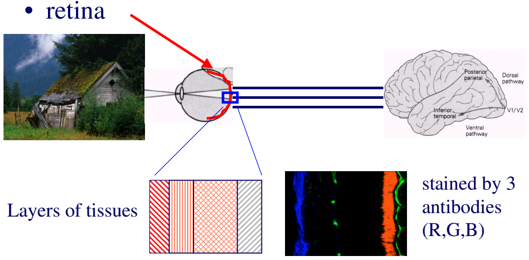

Computer Scientist's View of Retinal Detachment

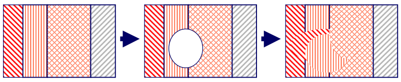

normal detachment 7 days after

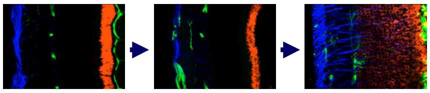

Detachment Development

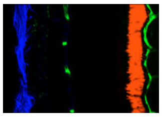

Normal

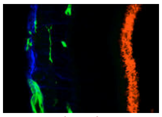

1 day after detachment

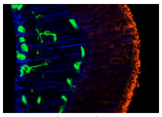

3 days after detachment

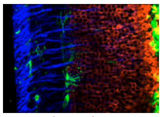

7 days after detachment

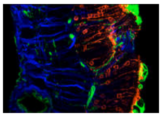

28 days after detachment

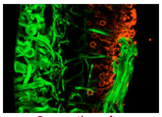

3 months after detachment

Data and Problem

- (Problem) What happens in retina after detachment?
  - What tissues (regions) are involved?
  - How do they change over time?
- **How will a program convey this info?**
- More than classification
  "we want to learn what classifier learned"

Visual vocabulary?

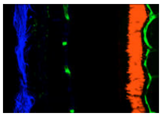

news:
president,
minister,
economic

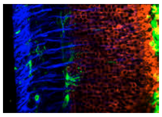

sports:
baseball,
score,
penalty

Visual Vocabulary (ViVo) generation

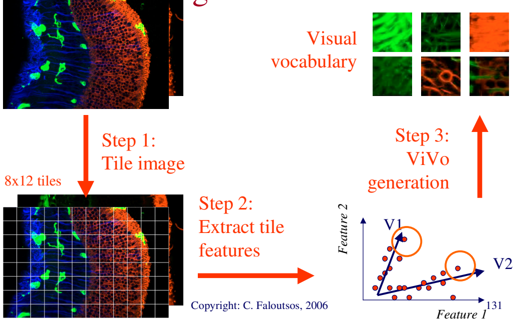

Biological interpretation

ID	ViVo	Description	Condition
V1	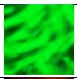	GFAP in inner retina (Müller cells)	Healthy
V10	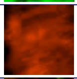	Healthy outer segments of rod photoreceptors	Healthy
V8	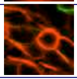	Redistribution of rod opsin into cell bodies of rod photoreceptors	Detached
V11	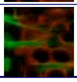	Co-occurring processes: Müller cell hypertrophy and rod opsin redistribution	Detached

Which tissue is significant on 7-day?

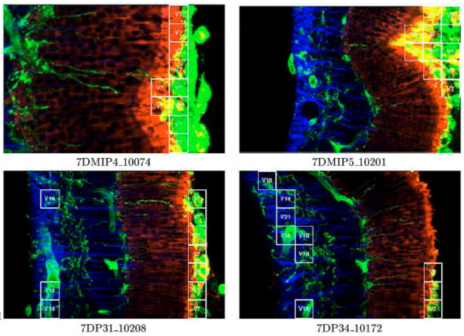

Automatic ViVo-annotation of images

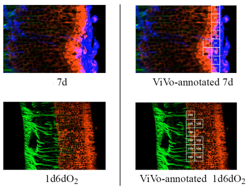

- A tile represents a ViVo v_k if the largest coefficient of the tile is along the k^{th} basis vector
- A ViVo v_k represents a class c_i if the majority of its tiles are in that class
- For each image, the representative ViVos for the class are automatically highlighted

Conclusion

- ICA can help us find ``good'' hidden variables (= visual vocabulary words)

Outline

- ...
- Tools and case studies: PCA, ICA, Random Walks
  - ICA
  - image captioning - Random Walks
- Conclusions

Image captioning: correlation between images and terms

Idea

- each modality into a set of nodes
- (could have many more modalities: audio, motion, etc etc)
- random walk, on the resulting m-partite graph

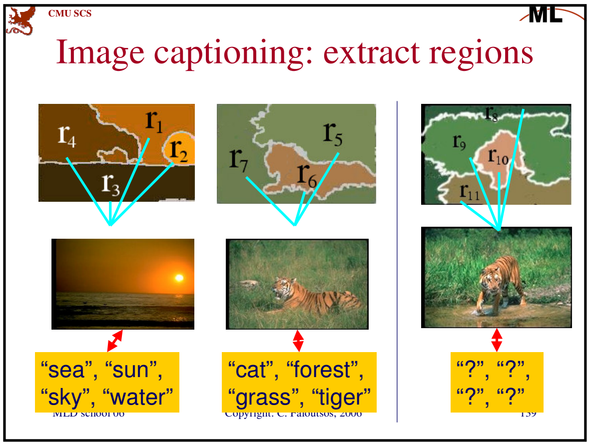
CMU SCS
ML
Image captioning: extract regions
“sea”, “sun”, “sky”, “water”
“cat”, “forest”, “grass”, “tiger”
“?”, “?”, “?”, “?”
MLD school 06
Copyright: C. Faloutsos, 2006
139

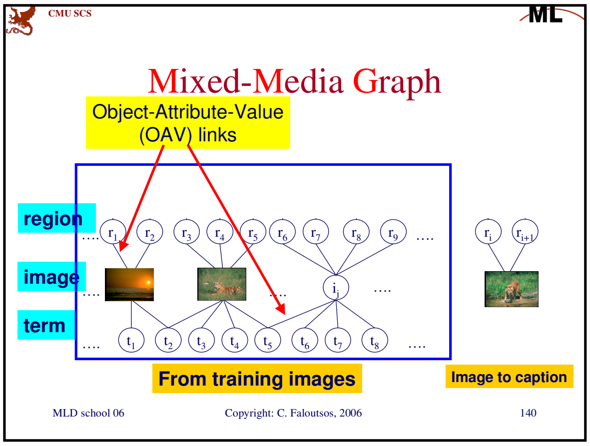
CMU SCS
ML
Mixed-Media Graph
Object-Attribute-Value (OAV) links
region
image
term
From training images
Image to caption
MLD school 06
Copyright: C. Faloutsos, 2006
140

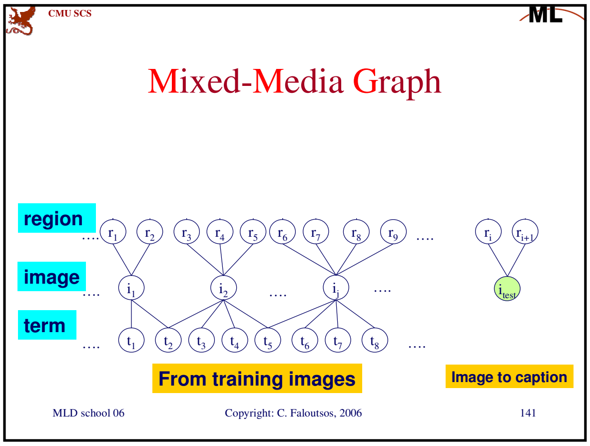
CMU SCS
ML
Mixed-Media Graph
region
image
term
From training images
Image to caption
MLD school 06
Copyright: C. Faloutsos, 2006
141

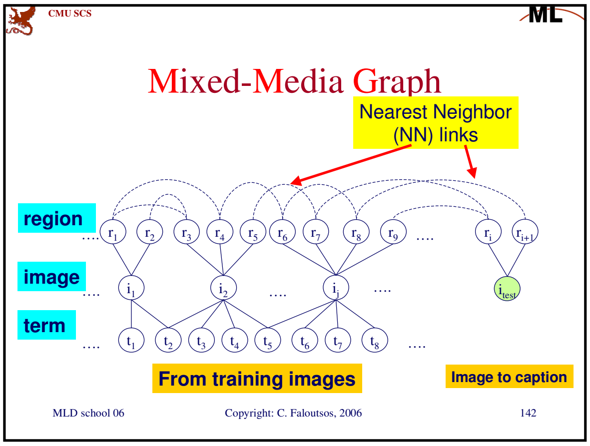
CMU SCS
ML
Mixed-Media Graph
Nearest Neighbor (NN) links
region
image
term
From training images
Image to caption
MLD school 06
Copyright: C. Faloutsos, 2006
142

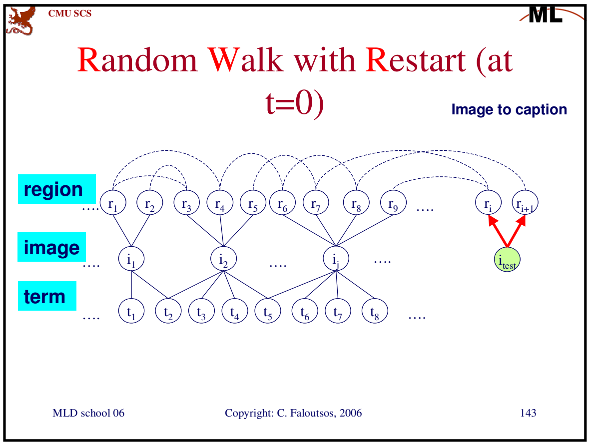
CMU SCS
ML
Random Walk with Restart (at t=0)
Image to caption
region
image
term
MLD school 06
Copyright: C. Faloutsos, 2006
143

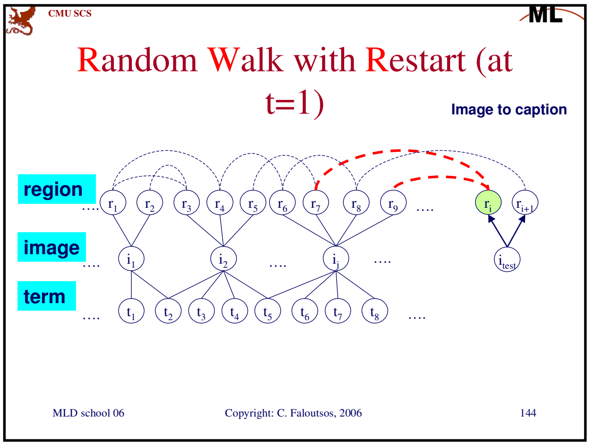
CMU SCS
ML
Random Walk with Restart (at t=1)
Image to caption
region
image
term
MLD school 06
Copyright: C. Faloutsos, 2006
144

CMU SCS

Random Walk with Restart (at t=2)

Image to caption

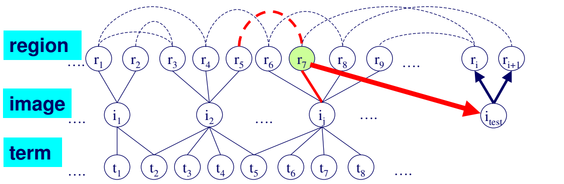

CMU SCS

Random Walk with Restart (at t=3)

Image to caption

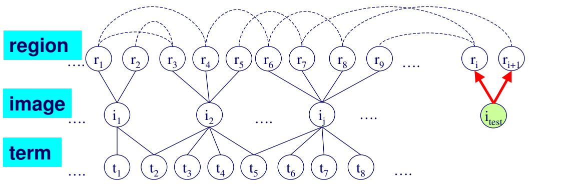

CMU SCS

Random Walk with Restart (at t=0)

Image to caption

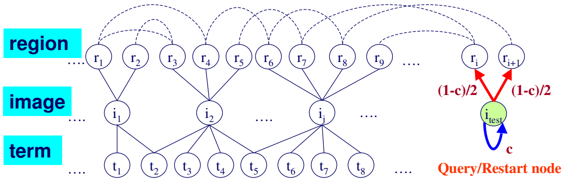

Blue step: restart at the query node (with probability **c**)
Red step: randomly walk among one link (with prob. **1-c**)

CMU SCS

Random Walk with Restart (at t>0)

Image to caption

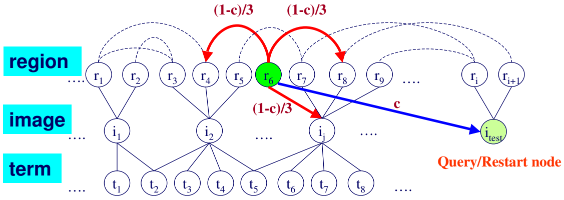

Blue step: restart at the query node (with probability **c**)
Red step: randomly walk among one link (with prob. **1-c**)

CMU SCS

Steady state probability of RWR

Image to caption

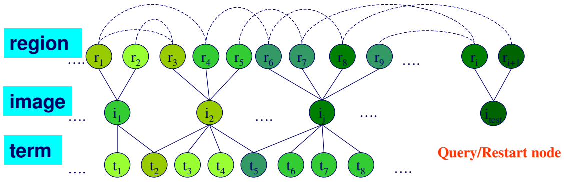

Dark green: higher probability to be visited
Light green: lower probability to be visited

CMU SCS

Steady state probability of RWR

Image to caption

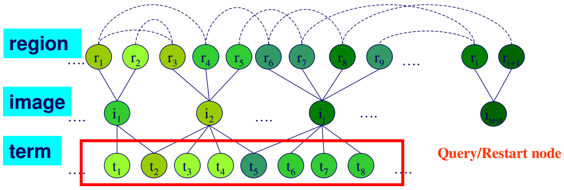

Dark green: higher probability to be visited
Light green: lower probability to be visited

Compute the steady state probability

$$\vec{\mathbf{u}} = (1-c)\mathbf{A}\vec{\mathbf{u}} + c\vec{\mathbf{v}}$$

$$\Rightarrow \vec{\mathbf{u}} = c(\mathbf{I} - (1-c)\mathbf{A})^{-1}\vec{\mathbf{v}}$$

$\vec{\mathbf{u}}$: steady state probability

$\vec{\mathbf{v}}$: restart vector

$\mathbf{A}$: transition matrix

c : restart probability

Putting MMG+RWR to work

- Tasks:
  - Image captioning: GCap [Pan+, MDDE 2004]
  - (Multi-modal retrieval: MMSS [Pan+, ICDM 2004])

GCap: MMG+RWR for image captioning

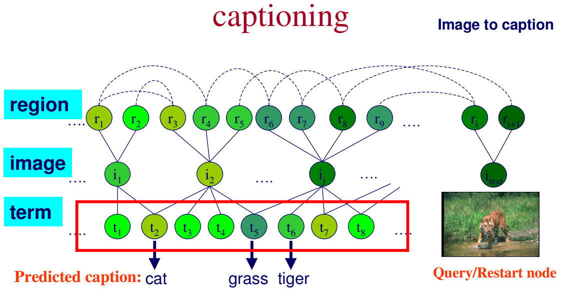

Examples of captioning results

Image			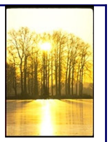
Truth	*cat, grass, tiger, water*	*mane, cat, lion, grass*	*sun, water, tree, sky*
Our caption	grass, cat, tiger, water	lion, grass, cat, mane	tree, water, building, sky

Predicted terms are listed in the order of likeliness.

Extension: other cross-modal queries

- e.g., text-to-text

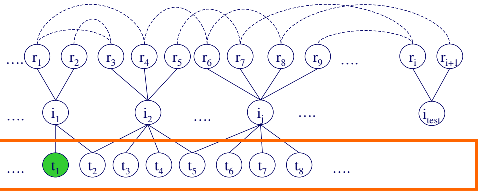

(Extension) Other cross-modal queries

- For example, the term-to-term correlation
  - "Given a term, find other similar terms."

Term	1	2	3	4	5
Branch	Birds	Night	Owl	Nest	Hawk
Bridge	Water	Arch	Sky	Stone	Boats
Car	Tracks	Street	Buildings	Turn	prototype

Extension: group captioning

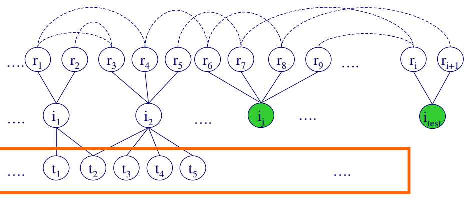

(Extension) Captioning images in groups

Image			
Truth	*sun, water, tree, sky*	*sun, clouds, sky, horizon*	*sun, water*
GCap caption	tree, people, sky, water	water, tree, people, sky	sky, sun
Group caption	sky, water, tree, sun		

Summary

- GCap = MMG+RWR, for image captioning
  - Easy to apply to all kinds of cross-modal data
  - Outperforms the best previous image captioning results

OVERALL CONCLUSIONS

Powerful tools for image/multimedia mining:

- GEMINI + R-trees for similarity search
- Feature extraction with
  - Wavelets,
  - PCA, ICA
- Cross-modal mining / auto-captioning:
  - graph-based methods + Random Walks

THANK YOU!

christos <at> cs.cmu.edu
www.cs.cmu.edu/~christos